Logic and Probabilistic Models of Belief Change

Eric Pacuit

Department of Philosophy
University of Maryland, College Park
pacuit.org

February 25, 2016

- Updating probabilities
- Joyce's accuracy-first justification of Probabilism
- How to proceed?
- Greaves and Wallace's justification of conditionalization

A basic result about probabilities.

For any finite partition $\{E_i\}$ of the state space and any event H,

$$p(H) = \sum_i p(H \mid E_i)$$

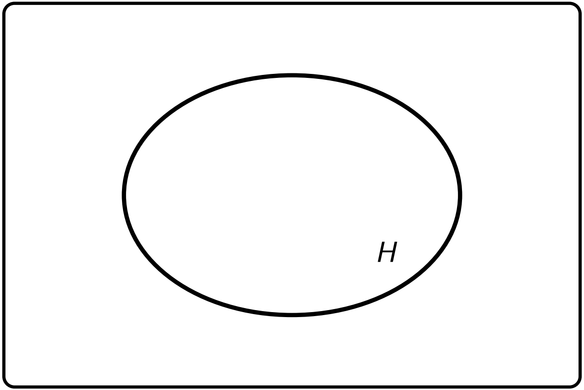

H

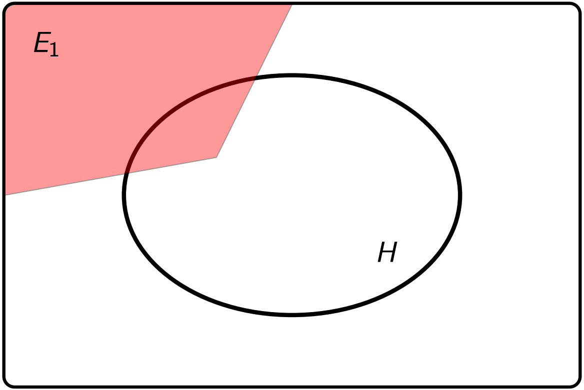
E_1
H

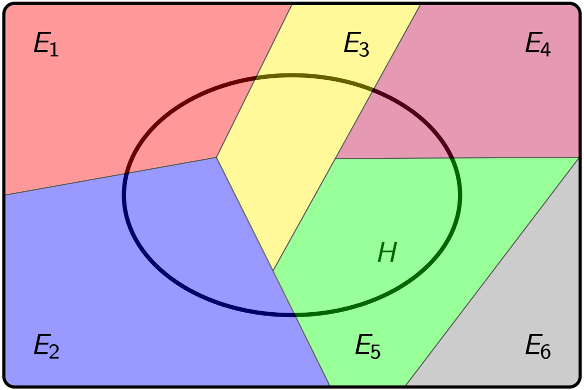
E_1
E_3
E_4
H
E_2
E_5
E_6

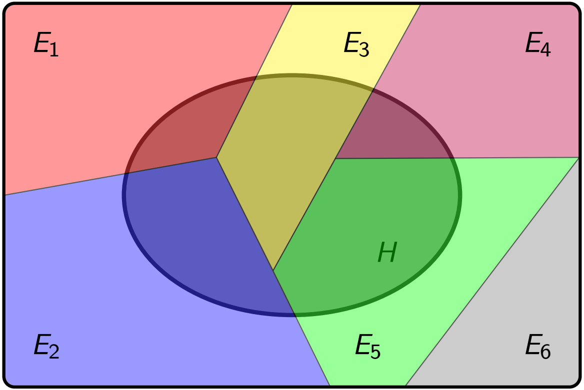

$$p(H) \quad = \quad p(H \cap E_1) + \cdots + p(H \cap E_6)$$

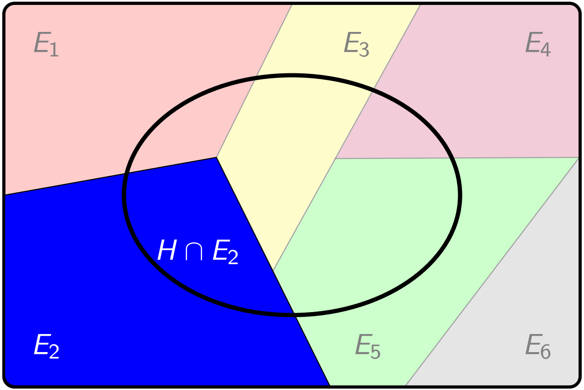

$$p(H) \quad = \quad p(H \cap E_1) + {\color{blue}p(H \cap E_2)} + \cdots + p(H \cap E_6)$$

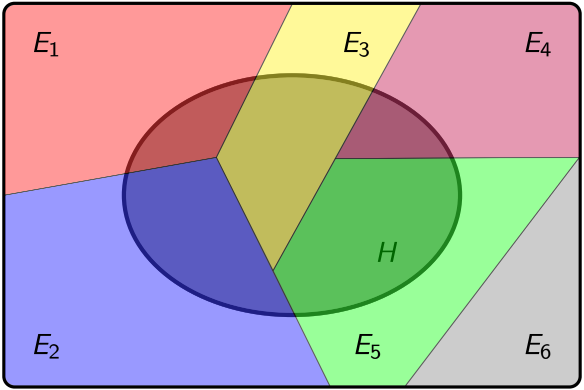

$$
\begin{aligned}
p(H) &= p(H \cap E_1) + \cdots + p(H \cap E_6) \\
&= \frac{p(E_1)}{p(E_1)} p(H \cap E_1) + \cdots + \frac{p(E_6)}{p(E_6)} p(H \cap E_6)
\end{aligned}
$$

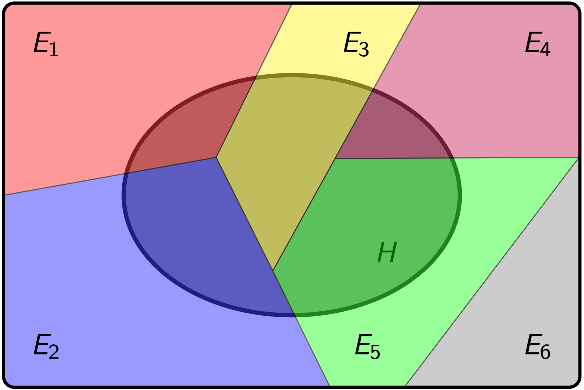

$$
\begin{aligned}
p(H) &= p(H \cap E_1) + \cdots + p(H \cap E_6) \\
&= \frac{p(E_1)}{p(E_1)} p(H \cap E_1) + \cdots + \frac{p(E_6)}{p(E_6)} p(H \cap E_6) \\
&= \sum_i p(E_i) P(H \mid E_i)
\end{aligned}
$$

Updating probabilities

Orthodox Bayesian Policy

- accept as admissible input only propositions;
- as response to such an input the only admissible change is conditioning the prior on the proposition in question.

Updating probabilities

Orthodox Bayesian Policy

- accept as admissible input only propositions;
- as response to such an input the only admissible change is conditioning the prior on the proposition in question.

Departing from a (orthodox) Bayesian policy:

1. accept as admissible a wider variety of inputs (e.g. expected values);
2. an admissible response to such an input can be a change in the prior that is not the result of conditioning;
3. an admissible response to such an input may be non-unique, that is, the posterior may not be uniquely determined by the prior + input.

p_0

E

$p_0(E) > 0$

$p(\cdot) = p_0(\cdot \mid E)$

$$p_0(\cdot, \top)$$

$$\downarrow$$

$$\boxed{E}$$

$$\downarrow$$

$$p(\cdot) = p_0(\cdot, E)$$

$$p_0$$

$$(E_1 : q_1, \ldots, E_k : q_k)$$
$$\{E_i\} \text{ is a partition, } \textstyle\sum_i q_i = 1$$

$$p(\cdot) = \textstyle\sum_i q_i * p_0(\cdot \mid E_i)$$

$$
\begin{array}{c}
p_0 \\
\downarrow \\
\boxed{\mathcal{C}\text{: set of constraints}} \\
\downarrow \\
p \text{ satisfies } \mathcal{C}
\end{array}
$$

MAXENT

Let us start with the simplest case, where our outcome space, X, contains only a finite number of points, $x_1, x_2, \ldots, x_n$. Then the **entropy** of a probability, p, on this space is:

$$-\sum_i p(x_i) \log p(x_i)$$

and the **information** is the negative of the entropy.

The minimum information or maximum entropy probability is the one which makes the states equiprobable: $p(x_i) = \frac{1}{n}$.

Consider three die x_1, x_2, x_3 and a random variable f such that $f(x_i) = i$.

$$\mathbb{E}[f] = p(x_1)f(x_1) + p(x_2)f(x_2) + p(x_3)f(x_3)$$

What probabilities maximize entropy under the constraint that $\mathbb{E}[f]$ have different values?

MAXENT

$\mathbb{E}[f]$	$p(x_1)$	$p(x_2)$	$p(x_3)$
1	1	0	0
0.1	0.907833	0.084333	0.007834
0.2	0.826297	0.147407	0.026297
$\vdots$	$\vdots$	$\vdots$	$\vdots$
0.8	0.438371	0.323257	0.238271
0.9	0.384586	0.330829	0.284586
2.0	0.333333	0.333333	0.333333
2.1	0.284586	0.330829	0.384586
2.2	0.238372	0.323257	0.438370
$\vdots$	$\vdots$	$\vdots$	$\vdots$
2.8	0.026297	0.147407	0.826296
2.9	0.007834	0.084332	0,907834
3.0	0	0	1

MAXENT

$\mathbb{E}[f]$	$p(x_1)$	$p(x_2)$	$p(x_3)$
1	1	0	0
0.1	0.907833	0.084333	0.007834
0.2	0.826297	0.147407	0.026297
$\vdots$	$\vdots$	$\vdots$	$\vdots$
0.8	0.438371	0.323257	0.238271
0.9	0.384586	0.330829	0.284586
2.0	0.333333	0.333333	0.333333
2.1	0.284586	0.330829	0.384586
2.2	0.238372	0.323257	0.438370
$\vdots$	$\vdots$	$\vdots$	$\vdots$
2.8	0.026297	0.147407	0.826296
2.9	0.007834	0.084332	0,907834
3.0	0	0	1

MAXENT

$\mathbb{E}[f]$	$p(x_1)$	$p(x_2)$	$p(x_3)$
1	1	0	0
0.1	0.907833	0.084333	0.007834
$\vdots$	$\vdots$	$\vdots$	$\vdots$
0.9	0.384586	0.330829	0.284586
2.0	0.333333	0.333333	0.333333
2.1	0.284586	0.330829	0.384586
$\vdots$	$\vdots$	$\vdots$	$\vdots$
2.9	0.007834	0.084332	0.907834
3.0	0	0	1

The MAXENT probabilities are **not** *closed under mxing*: A mixture of $(1, 0, 0)$ and $(0, 0, 1)$ is $(0.5, 0, 0.5)$, but this is not in the list...

Kullback-Leibler

Suppose that we start with a prior probability, p_0, and move to a posterior p_1 which satisfies certain constraints. The Kullback-Leibler "distance" is:

$$I(p_1, p_0) = \sum_i p_1(x_i) \log \frac{p_1(x_i)}{p_0(x_i)}$$

p_0

$\mathcal{C}$: set of constraints

p satisfies $\mathcal{C}$

p minimizes $I(p, p_0)$

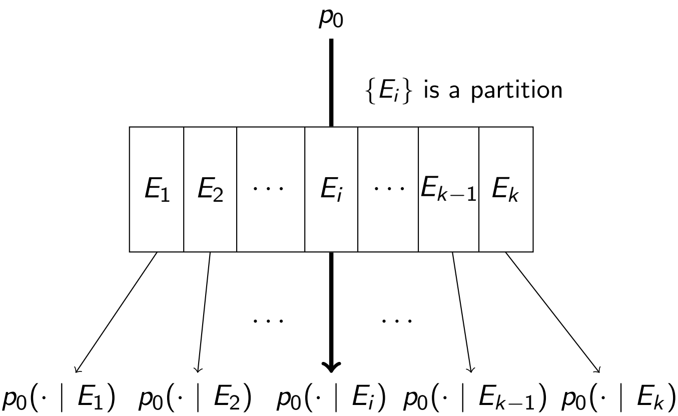
p_0
$\{E_i\}$ is a partition
E_1
E_2
$\cdots$
E_i
$\cdots$
E_{k-1}
E_k
$\cdots$
$\cdots$
$p_0(\cdot \mid E_1)$
$p_0(\cdot \mid E_2)$
$p_0(\cdot \mid E_i)$
$p_0(\cdot \mid E_{k-1})$
$p_0(\cdot \mid E_k)$

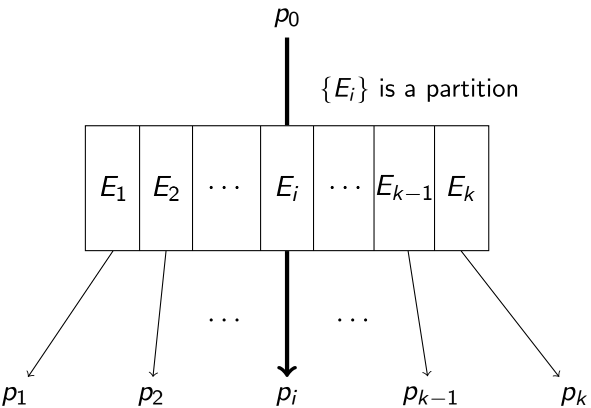
p_0
$\{E_i\}$ is a partition
E_1
E_2
$\cdots$
E_i
$\cdots$
E_{k-1}
E_k
$\cdots$
$\cdots$
p_1
p_2
p_i
p_{k-1}
p_k

Suppose that you are in a learning situation even more amorphous than the kind which motivates Jeffrey's idea. There is no nontrivial partition that you expect with probability one to be sufficient for your belief change....Perhaps you are in a novel situation where you expect the unexpected observational input....You are going to just think about some subject matter and update as a result of your thoughts...I will consider the learning situation a kind of black box and attempt no analysis of its internal structure.

(Skyrms, pg. 96, 97)

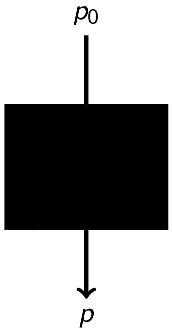
p_0
p

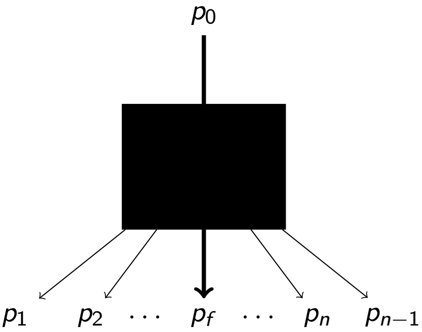
p_0
p_1 p_2 $\cdots$ p_f $\cdots$ p_n p_{n-1}

Martingale Property

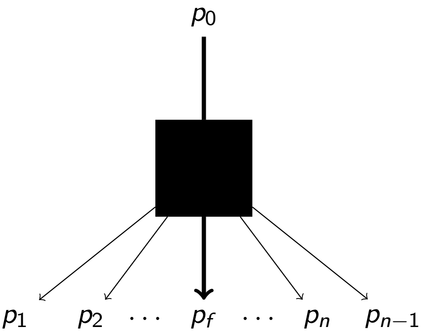

$$\text{(M)} \qquad p_0(A \mid p_f) = p_f(A)$$

It was suggested by Skyrms (1990) that this principle provides a plausible way to distinguish learning situations from situations where one expects probabilities to change for other reasons, such as getting drunk, having a brain lesion or having a dangerously low blood sugar level.

It was suggested by Skyrms (1990) that this principle provides a plausible way to distinguish learning situations from situations where one expects probabilities to change for other reasons, such as getting drunk, having a brain lesion or having a dangerously low blood sugar level.

[Huttegger] develops an account in which the reflection principle is a necessary condition for a black-box probability update to count as a *genuine learning experience*.

Simon Huttegger. *Learning Experiences and the Value of Knowledge*. Philosophical Studies, 2013.

The Value of Knowledge

The expected utility of an uninformed decision cannot be greater than the prior expectation of an informed decision.

The Value of Knowledge

Why is it better to make a “more informed” decision?
Suppose that you can either choose know, or perform a costless experiment and make the decision later. What should you do?

I. J. Good. *On the principle of total evidence*. British Journal for the Philosophy of Science, 17, pgs. 319 - 321, 1967.

“Never decide today what you might postpone until tomorrow in order to learn something new”

Choose between n acts $A_1, \ldots, A_n$ or perform a cost-free experiment E with possible results $\{e_k\}$, then decide.

$$EU(A) = \sum_i p(K_i)U(A \;\&\; K_i)$$

Choose between n acts $A_1, \ldots, A_n$ or perform a cost-free experiment E with possible results $\{e_k\}$, then decide.

$$EU(A) = \sum_i p(K_i)U(A \,\&\, K_i)$$

Then,

$$U(\text{Choose now}) = \max_j \sum_i p(K_i)U(A_j \,\&\, K_i)$$

$$= \max_j \sum_k \sum_i p(K_i)p(e_k \mid K_i)U(A_j \,\&\, K_i)$$

The value of an informed decision conditional on e:

$$\max_j \sum_i p(K_i \mid e) U(A_j \,\&\, K_i)$$

The value of an informed decision conditional on e:

$$\max_j \sum_i p(K_i \mid e)U(A_j \ \& \ K_i)$$

$$U(\text{Learn, Choose }) \quad = \quad \textstyle\sum_k p(e_k) \max_j \sum_i p(K_i \mid e_k)U(A_j \ \& \ K_i)$$

The value of an informed decision conditional on e:

$$\max_j \sum_i p(K_i \mid e)U(A_j \ \& \ K_i)$$

$$\begin{aligned} U(\text{Learn, Choose }) &= \textstyle\sum_k p(e_k)\max_j \sum_i p(K_i \mid e_k)U(A_j \ \& \ K_i) \\ &= \textstyle\sum_k p(e_k)\max_j \sum_i (\frac{p(e_k \mid K_i)p(K_i)}{p(e_k)})U(A_j \ \& \ K_i) \end{aligned}$$

The value of an informed decision conditional on e:

$$\max_j \sum_i p(K_i \mid e)U(A_j \ \& \ K_i)$$

$$\begin{aligned}
U(\text{Learn, Choose }) &= \textstyle\sum_k p(e_k) \max_j \sum_i p(K_i \mid e_k)U(A_j \ \& \ K_i) \\
&= \textstyle\sum_k p(e_k) \max_j \sum_i (\frac{p(e_k \mid K_i)p(K_i)}{p(e_k)})U(A_j \ \& \ K_i) \\
&= \textstyle\sum_k \max_j \sum_i p(e_k \mid K_i)p(K_i)U(A_j \ \& \ K_i)
\end{aligned}$$

The value of an informed decision conditional on e:

$$\max_j \sum_i p(K_i \mid e) U(A_j \,\&\, K_i)$$

$$\begin{aligned} U(\text{Learn, Choose }) &= \textstyle\sum_k p(e_k) \max_j \sum_i p(K_i \mid e_k) U(A_j \,\&\, K_i) \\ &= \textstyle\sum_k p(e_k) \max_j \sum_i (\frac{p(e_k \mid K_i) p(K_i)}{p(e_k)}) U(A_j \,\&\, K_i) \\ &= \textstyle\sum_k \max_j \sum_i p(e_k \mid K_i) p(K_i) U(A_j \,\&\, K_i) \end{aligned}$$

Compare $\max_j \sum_k \sum_i p(K_i) p(e_k \mid K_i) U(A_j \,\&\, K_i)$ and $\sum_k \max_j \sum_i p(e_k \mid K_i) p(K_i) U(A_j \,\&\, K_i)$

The value of an informed decision conditional on e:

$$\max_j \sum_i p(K_i \mid e)U(A_j \,\&\, K_i)$$

$$\begin{aligned}
U(\text{Learn, Choose }) &= \textstyle\sum_k p(e_k) \max_j \sum_i p(K_i \mid e_k)U(A_j \,\&\, K_i) \\
&= \textstyle\sum_k p(e_k) \max_j \sum_i (\frac{p(e_k \mid K_i)p(K_i)}{p(e_k)})U(A_j \,\&\, K_i) \\
&= \textstyle\sum_k \max_j \sum_i p(e_k \mid K_i)p(K_i)U(A_j \,\&\, K_i)
\end{aligned}$$

Compare $\max_j \sum_k \sum_i p(K_i)p(e_k \mid K_i)U(A_j \,\&\, K_i)$ and $\sum_k \max_j \sum_i p(e_k \mid K_i)p(K_i)U(A_j \,\&\, K_i)$

$\sum_k \max_j g(k, j)$ is greater than or equal to $\max_j \sum_k g(k, j)$, so the second is greater than or equal to the first.

(M) implies value of knowledge

Suppose that (M) holds.

(M) implies value of knowledge

Suppose that (M) holds.
Then (assuming that each $p(p_f)$ is positive) your value for choosing an act now is

(M) implies value of knowledge

Suppose that (M) holds.
Then (assuming that each $p(p_f)$ is positive) your value for choosing an act now is

$$\max_j \sum_i p(K_i)u(A_j \& K_i) = \max_j \sum_i \sum_f p(K_i \mid p_f)p(p_f)u(A_j \& K_i)$$

(M) implies value of knowledge

Suppose that (M) holds.
Then (assuming that each $p(p_f)$ is positive) your value for choosing an act now is

$$\max_j \sum_i p(K_i)u(A_j \& K_i) = \max_j \sum_i \sum_f p(K_i \mid p_f)p(p_f)u(A_j \& K_i)$$

$$= \max_j \sum_f \sum_i p_f(K_i)p(p_f)u(A_j \& K_i)$$

(M) implies value of knowledge

Suppose that (M) holds.
Then (assuming that each $p(p_f)$ is positive) your value for choosing an act now is

$$\max_j \sum_i p(K_i)u(A_j\&K_i) = \max_j \sum_i \sum_f p(K_i \mid p_f)p(p_f)u(A_j\&K_i)$$

$$= \max_j \sum_f \sum_i p_f(K_i)p(p_f)u(A_j\&K_i)$$

The value of choosing after the learning experience is:

$$\sum_f p(p_f) \max_j \sum_i p_f(K_i)u(A_j\&K_i)$$

(M) implies value of knowledge

Suppose that (M) holds.
Then (assuming that each $p(p_f)$ is positive) your value for choosing an act now is

$$\max_j \sum_i p(K_i)u(A_j \& K_i) = \max_j \sum_i \sum_f p(K_i \mid p_f)p(p_f)u(A_j \& K_i)$$

$$= \max_j \sum_f \sum_i p_f(K_i)p(p_f)u(A_j \& K_i)$$

The value of choosing after the learning experience is:

$$\sum_f p(p_f) \max_j \sum_i p_f(K_i)u(A_j \& K_i)$$

The latter term cannot be less than the former term on general mathematical grounds.

- The experiment is assumed to be essentially costless;
- You know that you are an expected utility maximizer and that you will be one after learning the true member of the partition.
- In the classical theorem you know that you will update by conditioning; in Skyrms' extension, you know that you will honor the martingale principle.
- By working within Savages decision theory, the states and acts are probabilistically independent (choosing an act does not give any information about the state).

- The states, acts and utilities are the same before and after the learning experience.

- Having the learning experience does not by itself alter your probabilities for states of the world (although the outcomes of the experience usually do); the learning experience and the states of the world are probabilistically independent.

...the martingale principle should not be applied to belief changes in epistemologically defective situations. In situations of memory loss, of being brainwashed or being under the influence of drugs, (M) should obviously not hold. If you believe that in an hour you will think you can fly because you're about to consume some funny looking pills, then you should not already now have that belief.

So, the martingale principle is claimed to apply if you learn something in the black-box, but not if you learn nothing or other things happen besides learning.

A genuine learning situation is partially characterized in the following way:

Postulate. If a belief change from p to $\{p_f\}$ constitutes a genuine learning situation, then

$$\sum_f p(p_f) \max_j \sum_i p_f(K_i)u(A_j \& K_i) \geq \max \sum_i p(K_i)u(A_j \& K_i)$$

for all utility values $u(A_j \& K_i)$ with strict inequality unless the same act maximizes expected utility irrespective of which of the p_f occurs.

If a belief change leads you to foreseeably make worse choices than you could already make now in some decision situations, then it cannot be a pure learning experience. Perhaps you are bolder after having taken those funny looking pills, for example. From your current perspective, this might help you in some decision problems, but it will be harmful in others.

F. Dietrich, C. List and R. Bradley. *Belief revision generalized: A joint characterization of Bayes's and Jeffrey's rules*. manuscript.

Characterization Result

Responsiveness: The agent's revised belief state *respects the constraint* given by the input.

Conservativeness: For all belief-input pairs (p, I), if I is "silent" on the probability of a (relevant) event A given another B, this conditional probability is preserved.

A decision-theoretic example

Ann, an employer, must decide whether to hire Bob, a job candidate. There is no time for a job interview, since a quick decision is needed. Ann is uncertain about whether Bob is competent or not; both possibilities have prior probability $\frac{1}{2}$. It would help Ann to know whether Bob has previous work experience, since this is positively correlated with competence, but gathering this information takes time.

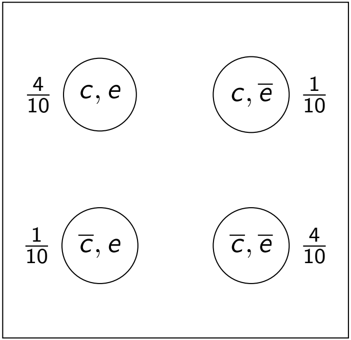

if $\gamma = c$, then $v(\gamma) = 5$

if $\gamma = \overline{c}$, then $v(\gamma) = -5$

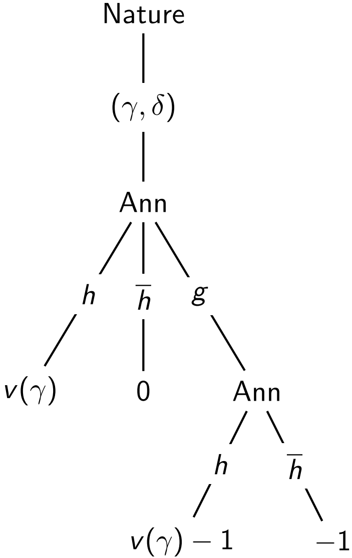

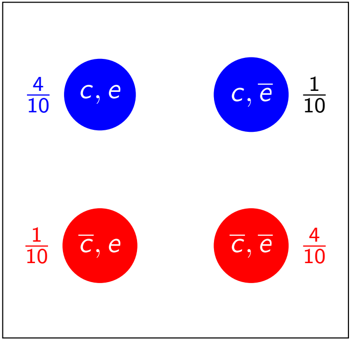

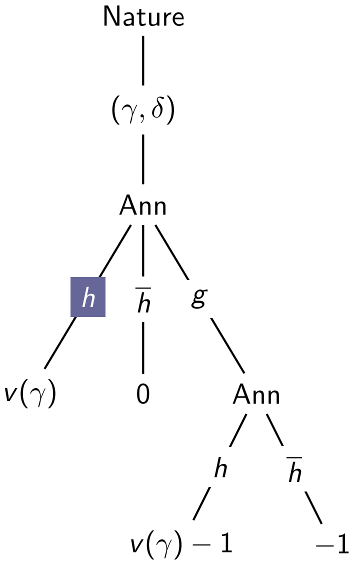

if $\gamma = c$, then $v(\gamma) = 5$

if $\gamma = \overline{c}$, then $v(\gamma) = -5$

$EU(h) =$

$p(C)u(C) + p(\overline{C})u(\overline{C}) =$

$0.5 * 5 + 0.5 * -5 = 0$

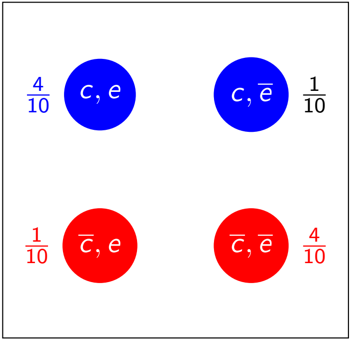

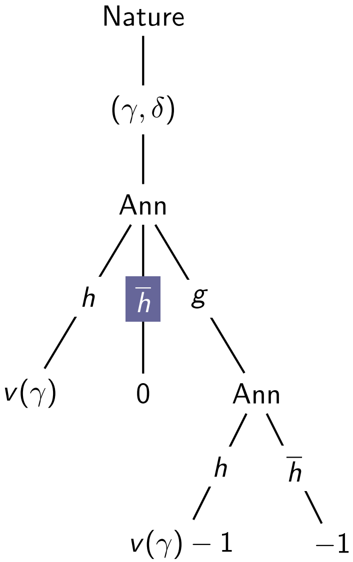

if $\gamma = c$, then $v(\gamma) = 5$

if $\gamma = \overline{c}$, then $v(\gamma) = -5$

$EU(\overline{h}) =$

$p(C)u(C) + p(\overline{C})u(\overline{C}) =$

$0.5 * 0 + 0.5 * 0 = 0$

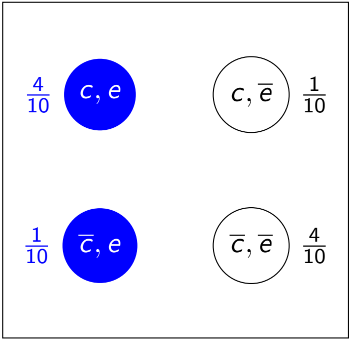

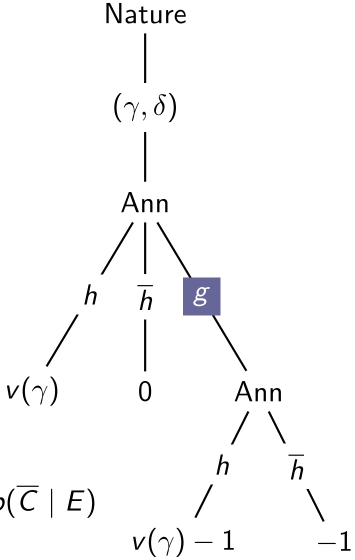

$$p(C \mid E) = \frac{\frac{4}{10}}{\frac{5}{10}} = \frac{4}{5}$$

$$(v(c) - 1)p(C \mid E) + (v(\overline{c}) - 1)p(\overline{C} \mid E)$$

$$EU_E(h) = 4\frac{4}{5} + -6\frac{1}{5} = 2$$

$$EU_E(\overline{h}) = -1$$

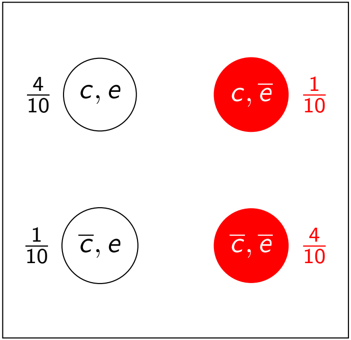

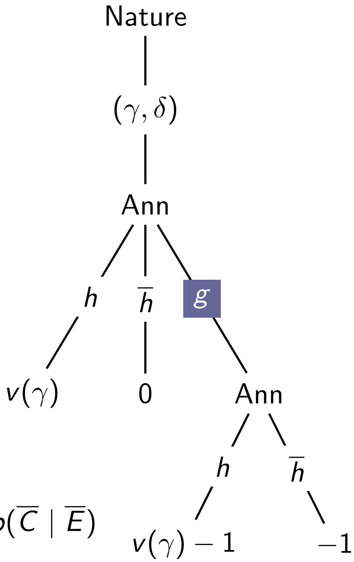

$$p(C \mid \overline{E}) = \frac{\frac{1}{10}}{\frac{5}{10}} = \frac{1}{5}$$

$$(v(c) - 1)p(C \mid \overline{E}) + (v(\overline{c}) - 1)p(\overline{C} \mid \overline{E})$$

$$EU_{\overline{E}}(h) = 4\frac{1}{5} + -6\frac{4}{5} = -4$$

$$EU_{\overline{E}}(\overline{h}) = 4\frac{1}{5} + -6\frac{4}{5} = -1$$

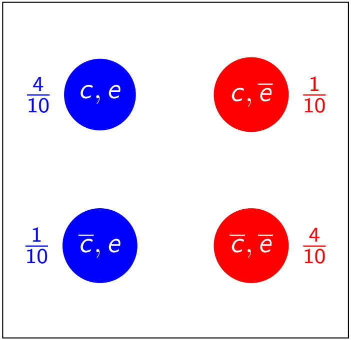

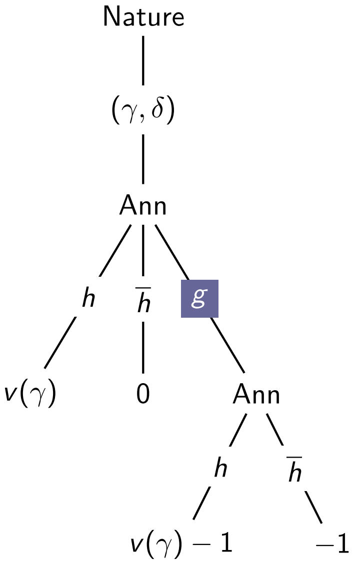

$$EU(g) =$$
$$p(E) * \max(EU_E(h), EU_E(\overline{h})) +$$
$$p(\overline{E}) * \max(EU_{\overline{E}}(h), EU_{\overline{E}}(\overline{h}))$$
$$= 0.5 * 2 + 0.5 * -1 = 0.5$$

Now suppose Ann follows her rational strategy. She writes to Bob to ask whether he has work experience. At this point, however, something surprising happens. Bob's answer reveals right from the beginning that his written English is poor. Ann notices this even before figuring out what Bob says about his work experience. In response to this unforeseen learnt input, Ann lowers her probability that Bob is competent from $\frac{1}{2}$ to $\frac{1}{8}$.

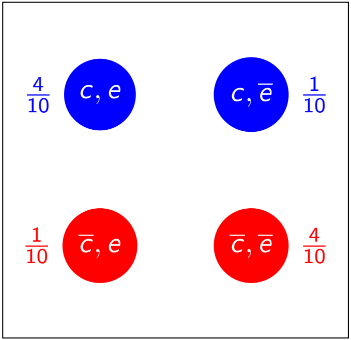

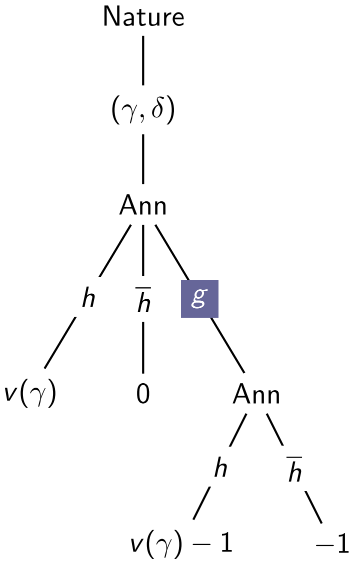

$$p_0 + (C : \tfrac{1}{8}, \overline{C} : \tfrac{7}{8}) = p_1$$

$$p_1(c, e) = \tfrac{4}{5} * \tfrac{1}{8} + 0 * \tfrac{7}{8} = \tfrac{1}{10}$$

$$p_1(c, \overline{e}) = \tfrac{1}{5} * \tfrac{1}{8} + 0 * \tfrac{7}{8} = \tfrac{1}{40}$$

$$p_1(\overline{c}, e) = 0 * \tfrac{1}{8} + \tfrac{1}{5} * \tfrac{7}{8} = \tfrac{7}{40}$$

$$p_1(\overline{c}, \overline{e}) = 0 * \tfrac{1}{8} + \tfrac{4}{5} * \tfrac{7}{8} = \tfrac{7}{10}$$

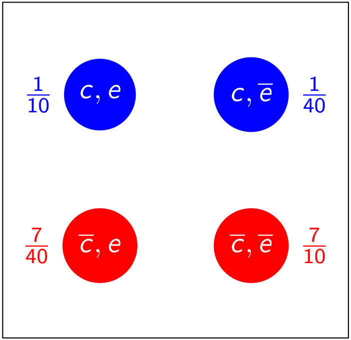

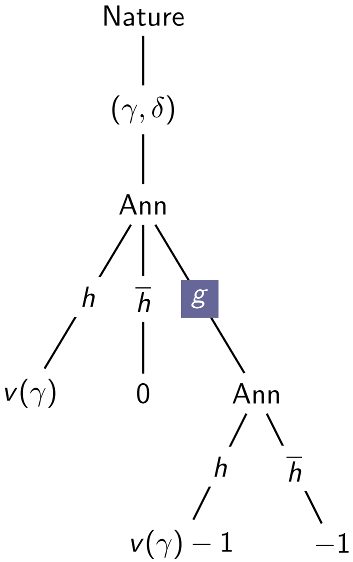

$$p_0 + (C : \tfrac{1}{8}, \overline{C} : \tfrac{7}{8}) = p_1$$

$$p_1(c, e) = \tfrac{4}{5} * \tfrac{1}{8} + 0 * \tfrac{7}{8} = \tfrac{1}{10}$$

$$p_1(c, \overline{e}) = \tfrac{1}{5} * \tfrac{1}{8} + 0 * \tfrac{7}{8} = \tfrac{1}{40}$$

$$p_1(\overline{c}, e) = 0 * \tfrac{1}{8} + \tfrac{1}{5} * \tfrac{7}{8} = \tfrac{7}{40}$$

$$p_1(\overline{c}, \overline{e}) = 0 * \tfrac{1}{8} + \tfrac{4}{5} * \tfrac{7}{8} = \tfrac{7}{10}$$

As she reads the rest of Bob's letter, Ann eventually learns that he has previous work experience, which prompts a Bayesian belief revision...

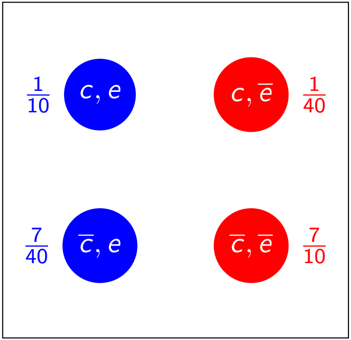

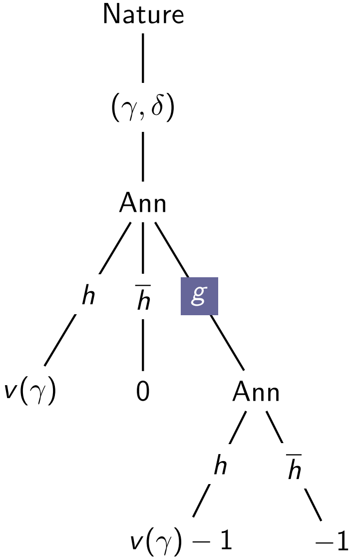

$EU_E(h) = 4\frac{4}{11} + -6\frac{7}{11} = -\frac{26}{11}$

$EU_E(\overline{h}) = -1$

Ann does not hire Bob.

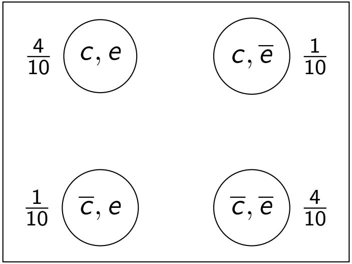

$\frac{4}{10}$
c, e
$c, \overline{e}$
$\frac{1}{10}$
$\frac{1}{10}$
$\overline{c}, e$
$\overline{c}, \overline{e}$
$\frac{4}{10}$

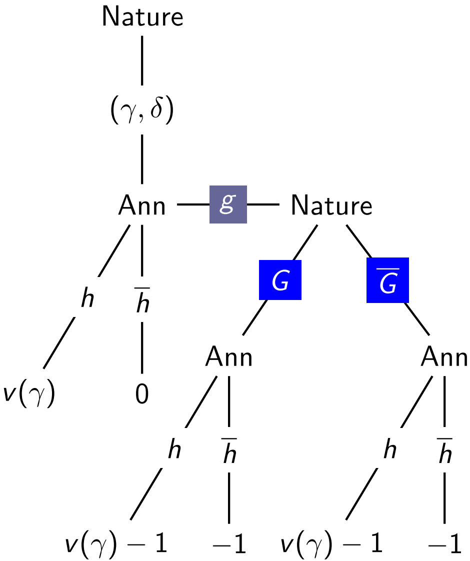

Nature
(γ, δ)
Ann
g
Nature
h
$\overline{h}$
G
$\overline{G}$
$v(\gamma)$
0
Ann
Ann
h
$\overline{h}$
h
$\overline{h}$
$v(\gamma) - 1$
-1
$v(\gamma) - 1$
-1

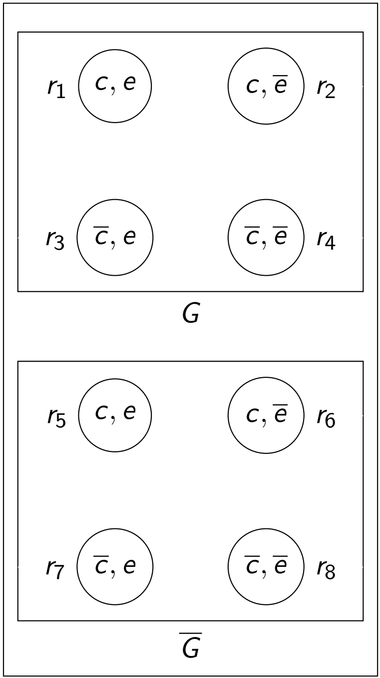
r_1
c, e
$c, \overline{e}$
r_2
r_3
$\overline{c}, e$
$\overline{c}, \overline{e}$
r_4
G
r_5
c, e
$c, \overline{e}$
r_6
r_7
$\overline{c}, e$
$\overline{c}, \overline{e}$
r_8
$\overline{G}$

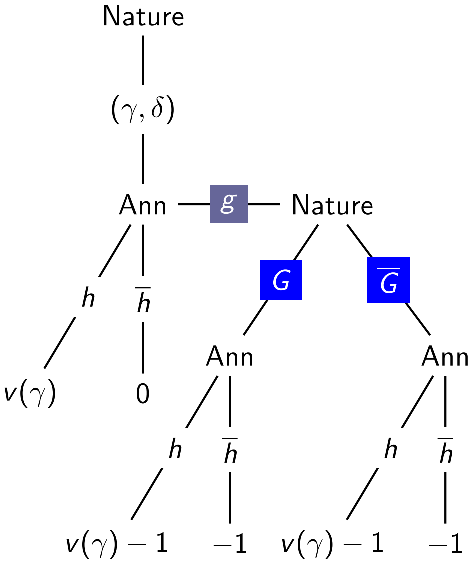
Nature
(γ, δ)
Ann
g
Nature
h
$\overline{h}$
G
$\overline{G}$
$v(\gamma)$
0
Ann
Ann
h
$\overline{h}$
h
$\overline{h}$
$v(\gamma) - 1$
-1
$v(\gamma) - 1$
-1

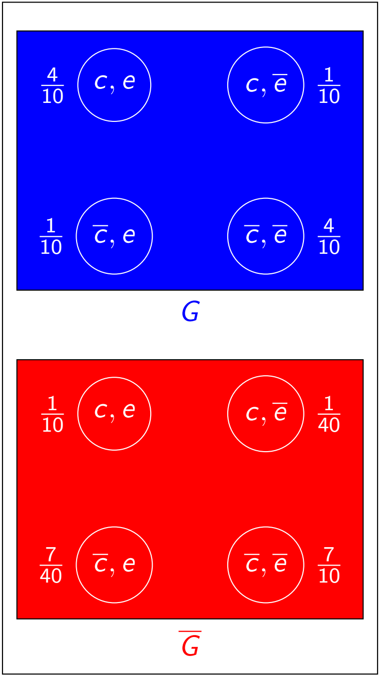
$\frac{4}{10}$ c, e
$c, \overline{e}$ $\frac{1}{10}$
$\frac{1}{10}$ $\overline{c}, e$
$\overline{c}, \overline{e}$ $\frac{4}{10}$
G
$\frac{1}{10}$ c, e
$c, \overline{e}$ $\frac{1}{40}$
$\frac{7}{40}$ $\overline{c}, e$
$\overline{c}, \overline{e}$ $\frac{7}{10}$
$\overline{G}$

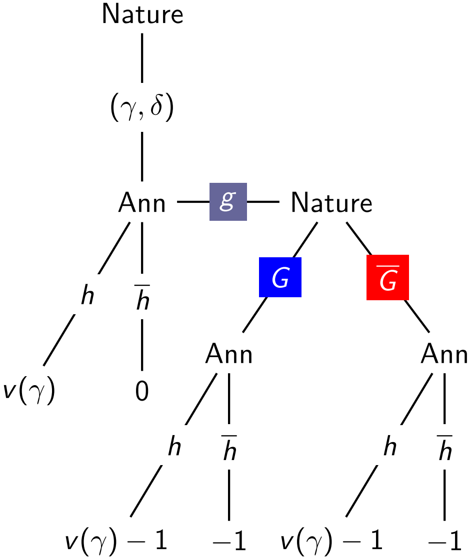
Nature
(γ, δ)
Ann
g
Nature
h
$\overline{h}$
G
$\overline{G}$
$v(\gamma)$
0
Ann
Ann
h
$\overline{h}$
h
$\overline{h}$
$v(\gamma) - 1$
-1
$v(\gamma) - 1$
-1

1. In her first decision (between h, $\overline{h}$, and g), Ann is falsely taken to foresee the possibilities of learning G or learning $\overline{G}$...This artificially complicates her expected-utility maximization exercise...

2. The additional conditionalization on G misrepresents Ann's beliefs, since the absence of linguistic errors in Bob's letter goes unnoticed.

3. Although it is true that the unforeseen news that Bob's written English is poor implies that Ann cannot uphold her original conceptualization of the decision problem, it does not follow that Ann re-conceptualizes her decision problem in line with the above model.

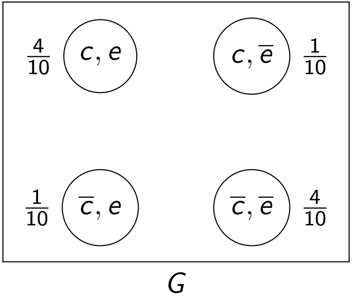

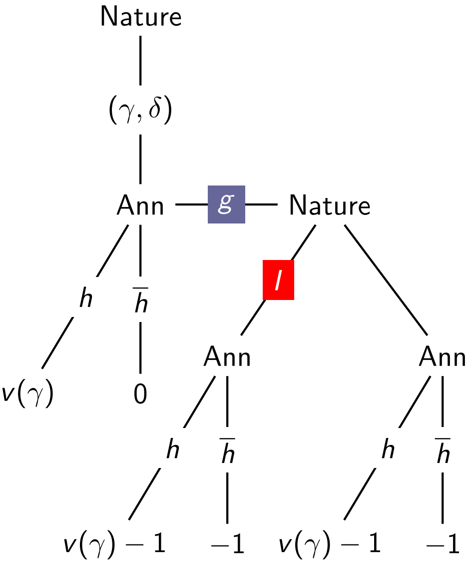

Nature does not reveal G or $\overline{G}$.

There is the possibility of a surprise move by Nature: Ann receives a particular unforeseen (Jeffrey) input I

Topics

- Classic papers (Makinson, Diaconis & Zabel, KLM, ...)
- Beliefs, credences and probability (Leitgeb's stability theory of belief, Pettigrew, Fitelson & Shear)
- Revising probabilities (List, Dietrich & Bradley, Halpern)
- Conditioning vs. learning (Osherson et al., Curpi et al., Skyrms)
- Context shifts (Halpern & Grünwald, Romeijn, Pettigrew)
- Lottery, Preface and Review paradox (Leitgeb, Easwaren & Fitelson)
- Iterated belief change, long-term dynamics, convergence results (Huttegger, EP)
- Bayesian reasoning, reasoning to the best explanation, case-base reasoning (Gilboa et al., Douven and Shubach)